

NASA/CR-1999-209331
ICASE Report No. 99-17



Plane Smoothers for Multiblock Grids: Computational Aspects

*Ignacio M. Llorente
Universidad Complutense, Madrid, Spain*

*Boris Diskin
ICASE, Hampton, Virginia*

*N. Duane Melson
NASA Langley Research Center, Hampton, Virginia*

*Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, VA
Operated by Universities Space Research Association*



National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

Prepared for Langley Research Center
under Contract NAS1-97046

May 1999

Available from the following:

NASA Center for Aerospace Information (CASI)
7121 Standard Drive
Hanover, MD 21076-1320
(301) 621-0390

National Technical Information Service (NTIS)
5285 Port Royal Road
Springfield, VA 22161-2171
(703) 487-4650

PLANE SMOOTHERS FOR MULTIBLOCK GRIDS: COMPUTATIONAL ASPECTS*

IGNACIO M. LLORENTE[†], BORIS DISKIN[‡], AND N. DUANE MELSON[§]

Abstract. Standard multigrid methods are not well suited for problems with anisotropic discrete operators, which can occur, for example, on grids that are stretched in order to resolve a boundary layer. One of the most efficient approaches to yield robust methods is the combination of standard coarsening with alternating-direction plane relaxation in the three dimensions. However, this approach may be difficult to implement in codes with multiblock structured grids because there may be no natural definition of global lines or planes. This inherent obstacle limits the range of an implicit smoother to only the portion of the computational domain in the current block. This report studies in detail, both numerically and analytically, the behavior of blockwise plane smoothers in order to provide guidance to engineers who use block-structured grids. The results obtained so far show alternating-direction plane smoothers to be very robust, even on multiblock grids. In common computational fluid dynamics multiblock simulations, where the number of subdomains crossed by the line of a strong anisotropy is low (up to four), textbook multigrid convergence rates can be obtained with a small overlap of cells between neighboring blocks.

Key words. robust multigrid methods, multiblock grids, anisotropic discrete operators

Subject classification. Applied Mathematics

1. Introduction and previous work. Standard multigrid techniques are efficient methods for solving many types of partial differential equations (PDE's), due to their optimal complexity (the required work is linearly proportional to the number of unknowns) [2], optimal memory requirements, and good efficiency and scalability in parallel implementations [10, 12]. Although highly efficient multigrid methods have been developed for a wide class of problems governed by PDE's, these methods are still underutilized in production and commercial codes [17]. One reason for this underutilization is that the textbook efficiency achieved in solving uniformly elliptic problems is not maintained in anisotropic problems; i.e., convergence rates of standard multigrid methods degrade on problems that have anisotropic discrete operators.

Several methods to deal with anisotropic operators have been studied in the multigrid literature. One popular approach is semicoarsening where the multigrid coarsening process is not applied uniformly to all of the coordinate directions [13, 14, 16, 19]. By selectively *not* coarsening the grid in certain directions, the anisotropy can be reduced on the coarser grids. This process makes the smoothing problem easier because only part of the high-frequency error (namely, the error components oscillating in the directions of coarsening) should be eliminated in relaxation.

Another approach for dealing with anisotropic problems is to develop robust smoothers that can eliminate all the high-frequency errors in the presence of strong anisotropies [11, 15, 18]. These robust smoothers can

*This research was supported by the National Aeronautics and Space Administration under NASA Contract No. NAS1-97046 while the first and second authors were in residence at the Institute for Computer Applications in Science and Engineering (ICASE), NASA Langley Research Center, Hampton, VA 23681-2199

[†]Departamento de Arquitectura de Computadores y Automática, Universidad Complutense, 28040 Madrid, Spain (email: llorente@dacya.ucm.es)

[‡]Institute for Computer Applications in Science and Engineering, Mail Stop 132C, NASA Langley Research Center, Hampton, VA 23681-2199 (email: bdiskin@icase.edu)

[§]Aerodynamic and Aeroacoustic Methods Branch, NASA Langley Research Center, Hampton, VA 23681-2199 (email: n.d.melson@larc.nasa.gov)

be efficiently implemented in the framework of a multigrid method with full coarsening. Plane-implicit relaxation schemes belong to this family of robust smoothers. These schemes are natural for structured grids. (It is more difficult to apply them to unstructured grids because there is no natural definition of a plane or even a line.) Other intermediate alternatives that combine implicit plane or line relaxations with partial and full coarsening have also been presented in the multigrid literature [6, 8, 21].

Previously [11], we studied the behavior of plane-relaxation methods as multigrid smoothers on single-block grids in the three dimensions (3-D). With numerical experiments and the local mode analysis, we compared different methods by considering the dependence of their smoothing factors on the anisotropy strength. It was observed that for strong anisotropies and practical grid sizes, there are significant differences in the smoothing factors, depending on whether periodic or Dirichlet boundary conditions are employed. For Dirichlet boundary conditions, increasing anisotropy turns the smoother into an exact solver. (The convergence factor is inversely proportional to the anisotropy strength.)

It has long been known (see, e.g., [2]) that the alternating-direction plane relaxation scheme is a simple, extremely efficient and robust method. Its efficiency, in fact, is improved in the presence of strong anisotropies. This scheme was considered to be essentially more expensive (in work-count per relaxation sweep) than point relaxation schemes. However, it was found (see [11]) that the overall convergence rate of the 3-D multigrid solver does not deteriorate if, instead of exactly solving two-dimensional (2-D) problems arising in plane-relaxation sweeps, just one 2-D V-cycle is used in each plane. With this improvement, the convergence factor per work unit of the 3-D V-cycle employing the alternating-direction plane relaxation scheme for an *isotropic* problem is just twice that of the most efficient cycle using a pointwise smoother. This cost does not look excessive, taking into account the excellent robustness of the alternating-direction plane relaxation smoother in treating anisotropic problems.

Such single-block algorithms are very efficient (especially for structured grids with stretching) due to their relatively easy implementation (on both sequential and parallel computers) and good architectural properties (parallelism and cache memory exploitation). However, multiblock grids are needed to deal with complex geometries or to facilitate parallel processing. This multiblock approach is already useful in serial computers: it improves the data locality properties and efficiently exploits the memory hierarchy of the underlying computer [7].

Plane smoothers have been successfully applied to solve the incompressible Navier-Stokes equations on *rectangular* multiblock grids (multiblock grids where lines and planes are globally defined) [15]. In a general case, however, the plane-implicit smoother can be applied only within the current block, and so this plane smoother becomes a blockwise plane smoother. The purpose of the current work is to study whether the optimal properties of plane-implicit smoothers deteriorate for general multiblock grids where these smoothers are not applied globally, but inside each block.

Previous analysis for the 2-D case reported by Jones and Melson [9] suggests that multiblock grids have a detrimental effect on the performance of implicit schemes. By using numerical experiments and rigorous analysis, Jones and Melson derived relations between the block size, the amount of overlap between blocks, and the strength of the anisotropy, that must hold for the resulting multigrid algorithm to be efficient. By looking at the model problem, they showed that textbook multigrid efficiency is achieved only when the block sizes (and, therefore, the range of the implicit operators) are proportional to the strength of the anisotropy.

We have developed a flexible 3-D code to study the behavior of blockwise plane smoothers and determine whether the result relating the convergence rate with the block size, the overlap size, and the anisotropy strength obtained for the 2-D case holds for 3-D. The current version of the code solves the nonlinear

diffusion-convection equation by using a full multigrid approach with the *full approximation scheme* [2, 20] for V-cycles. This report presents results for the linear anisotropic diffusion equation. This equation can be solved in a multiblock, cell-centered, stretched grid for complex geometries. Each block can overlap with neighboring blocks. The code is implemented in Fortran77 and has been parallelized with the standard OpenMP directives for shared-memory parallel computing.

The research code is described in section 2. We have analyzed two cases of anisotropy. The first case is an anisotropic equation discretized on uniform grids (section 3) and the second case is an isotropic equation discretized on stretched grids (section 4). Different strategies for parallelizing multigrid solvers with blockwise relaxation schemes are outlined in section 5. Finally, section 6 presents some conclusions and future research directions.

2. The numerical problem. We will study the behavior of plane smoothers on multiblock grids when solving the anisotropic diffusion equation

$$(2.1) \quad a \frac{\partial^2 u(x, y, z)}{\partial x^2} + b \frac{\partial^2 u(x, y, z)}{\partial y^2} + c \frac{\partial^2 u(x, y, z)}{\partial z^2} = f(x, y, z)$$

on a 3-D rectangular open domain, Ω , with suitable boundary conditions on $\delta\Omega$, where u is an unknown function, and f is a specified source function. Coefficients a, b , and c in equation (2.1) can, in general, be functions of the spatial variables. A finite volume discretization of equation (2.1) is applied on a cell-centered computational grid Ω^1 . The discretization is given by the system of equations $L^1 u^1 = f^1$.

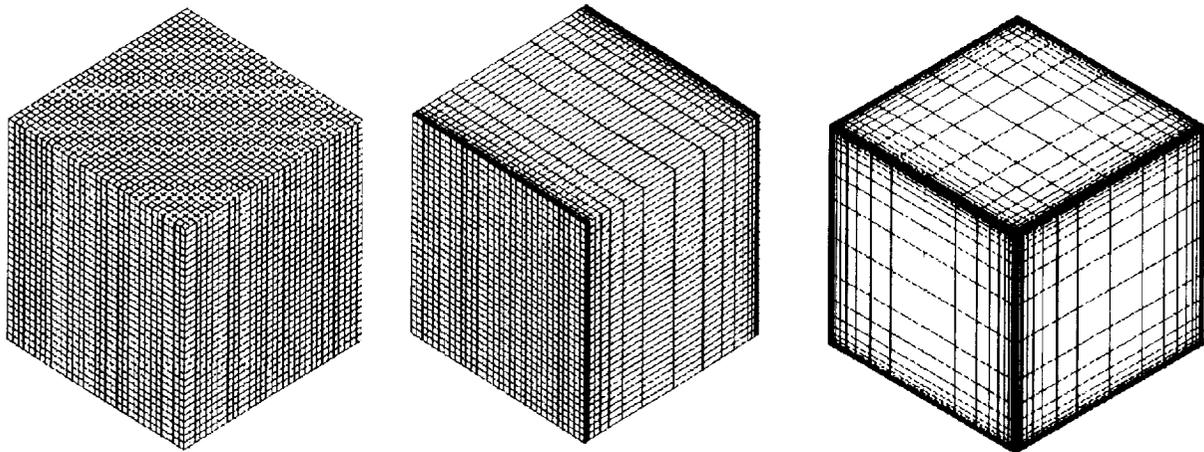


FIG. 1. $32 \times 32 \times 32$ uniform grid, $32 \times 32 \times 32$ grid stretched along x -direction (stretching ratio 1.5) and $32 \times 32 \times 32$ grid stretched along all directions (stretching ratio 1.5).

This paper studies two sources of anisotropy: anisotropic equation coefficients and nonunitary cell aspect ratios due to grid stretching. These sources reflect different situations: the former represents problems with a uniform anisotropy throughout the domain, while the latter exhibits problems with an anisotropy that varies from cell to cell. Grid stretching is commonly used in computational fluid dynamics (CFD) grid generation to pack points in regions with large solution gradients while avoiding an excess of points in more benign regions. In the present work, the stretching of the grid in a given direction is characterized by the stretching ratio (quotient between two consecutive mesh sizes in the same direction). See Figure 1 for examples of stretched grids. Varying local grid aspect ratios cause the mesh sizes (and, therefore, the anisotropy strength) to be

different in each cell. Note that for an exponential stretching in all the directions, every cell can have a different anisotropy.

2.1. The single-block algorithm. Our research code implements the *full approximation scheme* (FAS) to deal with nonlinearity. However, in the simplified (linear problem) case studied in this report the FAS performance is exactly the same as for the *correction scheme*.

A sequence of grids $\Omega^l (l = 1, \dots, M)$ is used in the multigrid scheme where Ω^1 is the finest target grid and the rest of the grids are obtained by applying cell-centered coarsening. In nonstretched-grid solvers, $h_l = 2h_{l-1}$ and $h_1 = h$. $L^l u^l = f^l$ is the discretization of equation (2.1) on Ω^l . The following iterative algorithm represents a $V(\gamma_1, \gamma_2)$ -cycle to solve the system $L^1 u^1 = f^1$.

step 1: Apply γ_1 sweeps of the smoothing method to $L^1 u^1 = f^1$

RESTRICTION PART

for $l = 2$ to M

step 2: Compute the residual $r^{l-1} = f^{l-1} - L^{l-1} u^{l-1}$

step 3: Restrict the residual $r^l = I_l^{l-1} r^{l-1}$

step 4: Restrict the current approximation $v^l = \hat{I}_l^{l-1} u^{l-1}$

step 5: Compute the right-hand side function $f^l = r^l + L^l v^l$

If $(l < M)$ then

step 6: Apply γ_1 sweeps of the smoothing method to $L^l u^l = f^l$

else

step 7: Solve the problem $L^M u^M = f^M$ on the coarsest grid

PROLONGATION PART

for $l = M - 1$ to 1

step 8: Correct the current approximation $u^l = u^l + I_{l+1}^l (u^{l+1} - v^{l+1})$

step 9: Apply γ_2 sweeps of the smoothing method to $L^l u^l = f^l$

The intergrid operators that perform the restriction (I_l^{l-1} at step 3 and \hat{I}_l^{l-1} at step 4) and the prolongation (I_{l+1}^l at step 8) connect the grid levels. The prolongation operator maps data from the coarser level to the current one while the restriction operators transfer values from the finer level to the current one. For the restriction operators, we used unweighted averaging for the solution u and a volume-weighted sum for the residual r . The prolongation operator is trilinear interpolation in the computational space.

Alternating-direction plane smoothers (Figure 2) for steps 1, 6, and 9 (in combination with full coarsening) have been found to be highly efficient and robust for anisotropic discrete operators on single-block grids [11]:

- *Optimal work per cycle.* An exact solution of each plane is not necessary; an approximate solution obtained by a single, 2-D multigrid cycle gives the same convergence rate of the 3-D multigrid cycle as an exact solution of each plane, but in much less execution time.
- *Very low convergence factor.* The convergence rate improves as the anisotropy becomes stronger.

2.2. Multiblock grids. Multiblock grids divide domain Ω into P subdomains Ω_p ($p = 1, \dots, P$)

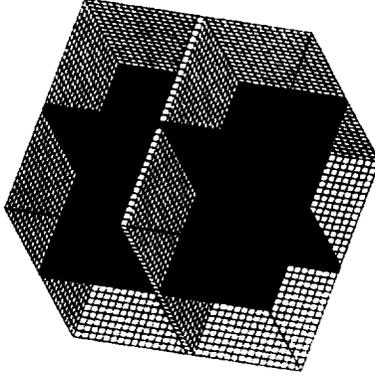


FIG. 2. *Alternating-direction plane-implicit smoother.*

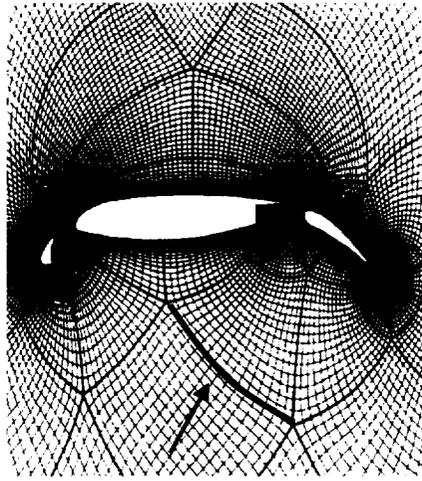


FIG. 3. *Multiblock grid for a multielement airfoil. Neighboring subgrids share a grid plane.*

$$(2.2) \quad \Omega = \cup_{p=1}^P \Omega_p,$$

where each subdomain Ω_p is covered with a structured grid Ω_p^1 . The grid blocking is generated by the grid generation software to deal with geometric complexities (Figure 3). A similar blocking is induced on the coarser grid levels. For multigrid, we construct a sequence of grids Ω^l ($l = 1, \dots, M$); each grid is defined as the union of P blocks:

$$(2.3) \quad \Omega^l = \cup_{p=1}^P \Omega_p^l; \quad (l = 1, \dots, M).$$

We assume that the grid lines are contiguous (possessing C^0 continuity) across the block interfaces (Figure 3).

This decomposition generates artificial boundaries within the original domain. It is necessary to apply some boundary conditions to the governing equations at these interfaces. In practice, a Dirichlet boundary condition is applied at artificial boundary cells (Figure 4). (Artificial boundary cells are cells adjacent to the

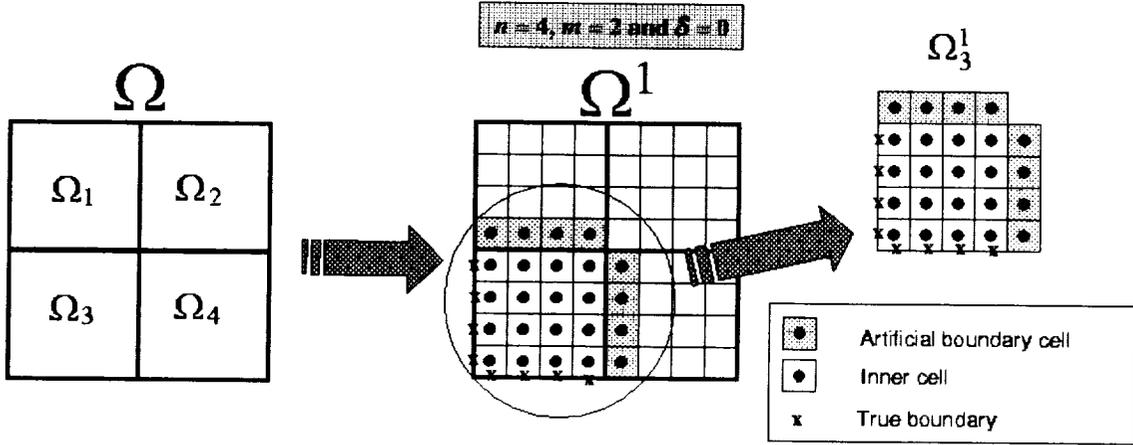


FIG. 4. Data structure of a subgrid without overlap; n is the number of cells per side in every block, m is the number of blocks per side in the split, and δ is the number of overlapping cell planes.

boundary of the given block.) The values of the unknowns assigned to these boundary cells are derived from the values at the corresponding cells in the neighboring blocks. Therefore, subgrids are extended to include these artificial boundary cells.

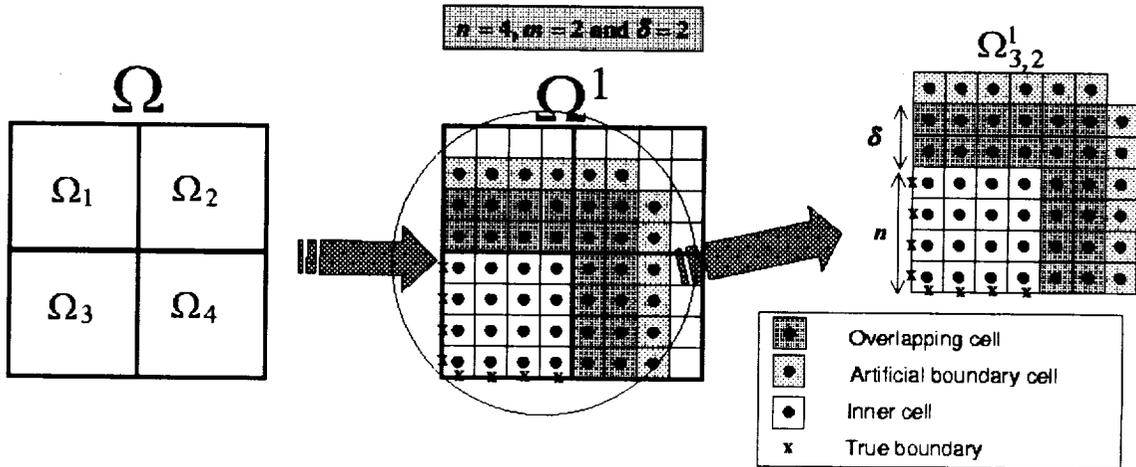


FIG. 5. Data structure of a subgrid with overlap: n is the number of cells per side in every block, m is the number of blocks per side in the split, and δ is the number of overlapping cell planes.

2.3. Overlapping subdomains. Let us define the extended subgrid $\Omega_{p,\delta}^l$ which is the subgrid Ω_p^l including all the inner and artificial boundary cells, plus an external margin of δ -cells deep overlapping into the neighboring block (Figure 5). In the general case of a cubic (3-D) partitioning when the overlap can be different in each of the six directions, δ is a six-component vector $(\delta_x^p, \delta_x^m, \delta_y^p, \delta_y^m, \delta_z^p, \delta_z^m)$. The expressions $L_p^l u_p^l = f_p^l$ and $L_{p,\delta}^l u_{p,\delta}^l = f_{p,\delta}^l$ denote discretizations on Ω_p^l and $\Omega_{p,\delta}^l$, respectively.

For simplicity, we consider a uniform overlapping (the same δ) throughout all the block interfaces. We must distinguish between the number of overlapping cells (2δ) and the physical size of the overlap $((2\delta+1)h)$. The latter is defined as the distance between centers of the artificial boundary cells of the two extended subgrids sharing a block interface. There are two strategies for the coarse-level block overlapping:

- *Fixed overlap.* The overlap parameter δ is the same in all the levels of the multigrid hierarchy

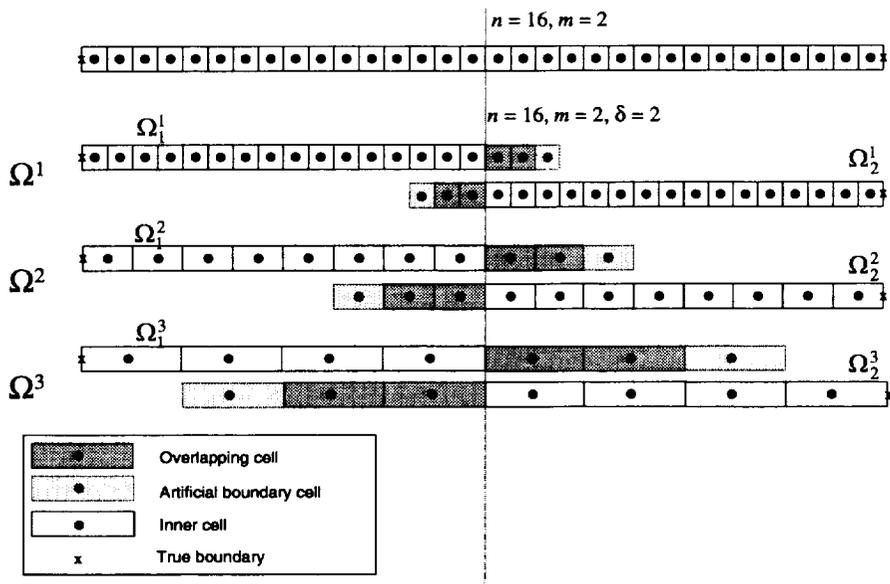


FIG. 6. Grid hierarchy of overlapping subdomains with a fixed number of overlapping cells. n is the number of cells per side in every block, m is the number of blocks per side in the split, and δ is the number of overlapping cell planes.

(Figure 6). Consequently, the overlapping space increases on the coarse levels. (It is doubled in each coarsening step.) This expansion happens because the number of the overlapping cells in the cross-interface direction remains fixed, while the size of each cell is doubled.

- *Decreasing overlap.* The number of overlapping cells in the cross-interface direction is reduced (δ is divided by 2 for each coarser level; see Figure 7). The overlapping space may also increase but much more slowly than in the previous case.

Again, in order to compare these strategies, we have to consider the performance (i.e., convergence rate per work unit) of the obtained algorithms. The work-count per cycle is a bit higher in the algorithm with a fixed overlap, due to the memory and computation overhead on the coarse levels. However, these two alternatives do not exhibit the same convergence rate (see discussion in section 3). Much of the efficiency of the multigrid algorithm is lost when using a decreasing overlap. Consequently, the fixed-overlap strategy was found to be globally more efficient.

The grid generation code generates a multiblock grid decomposed into subgrids. Extended subgrids are defined at the beginning of the simulation code. Note that overlapping cells may belong to several (more than 2) subgrids. This inherently implies a deterioration of the efficiency in a parallel setting.

In a general multiblock grid, there is no global definition of a line or a plane and so we have two alternatives to implement a V-cycle with an implicit smoother:

- *Domain decomposition with multigrid* applies the V-cycle inside each block.
- *Multigrid with blockwise plane smoother (Multigrid with domain decomposition smoother)* applies the V-cycle on the entire multiblock domain, while the smoothers (steps 1, 6, and 9) are performed

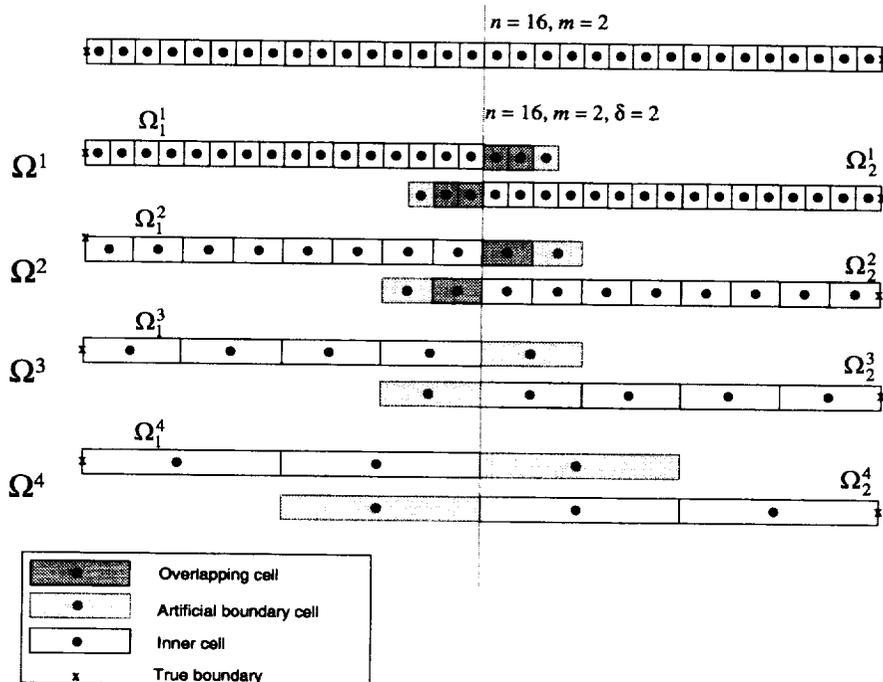


FIG. 7. Grid hierarchy of overlapping subdomains with a decreasing number of overlapping cells; n is the number of cells per side in every block, m is the number of blocks per side in the split, and δ is the number of overlapping cell planes.

inside each block.

2.4. Domain decomposition technique with multigrid. The following iterative algorithm represents a domain decomposition $V_{DD}(\gamma_1, \gamma_2)$ -cycle to solve the system $L^1 u^1 = f^1$, where Ω^1 is a multiblock grid decomposed into P subgrids.

for $p = 1$ to P

step 1: Apply $V(\gamma_1, \gamma_2)$ -cycle to solve the system $L_{p,\delta}^1 u_{p,\delta}^1 = f_{p,\delta}^1$ on $\Omega_{p,\delta}^1$

update 1: Update the overlap cells of neighboring blocks with values obtained in step 1

This technique is widely known by the computational community; it is the one-level multiplicative approach to solve the problem in a domain decomposition manner [1]. This alternative does not present a high level of parallelism because blocks are updated in sequential order. Other alternatives (Jacobi or colored ordering) with a higher parallelism can be found in the domain decomposition literature. The main drawback of the one-level multiplicative approach is its slow convergence rate. The convergence of this technique cannot be better than the convergence of the Schwartz method, which is known to be very poor. In fact, the convergence of a multigrid cycle with blockwise plane smoother will be bounded above by the convergence rate of the domain decomposition solver.

2.5. Multiblock with blockwise smoother. This approach is a priori more efficient because the multigrid algorithm is applied on the whole domain. Its efficiency and excellent parallelization potential

already have been demonstrated for *isotropic* elliptic problems [4] where pointwise red-black smoothers were used. In the proposed multigrid method for multiblock grids, the plane smoothers are applied within each block because no global definition of plane or line is assumed. However, the rest of the operators (restriction and prolongation) are completely local and do not need any global information. This method is a simple modification of the one-block approach, with the smoother applied blockwise. The following iterative algorithm represents a V_{BW} -cycle with the blockwise smoother to solve the system $L^1 u^1 = f^1$, where Ω^1 is a multiblock grid.

Apply γ_1 sweeps of the blockwise smoother:

for $p = 1$ to P

step 1: Apply the smoothing method to the system $L_{p,\delta}^1 u_{p,\delta}^1 = f_{p,\delta}^1$

update 1: Update the overlap cells of neighboring blocks with $u_{p,\delta}^1$

RESTRICTION PART

for $l = 2$ to M

for $p = 1$ to P

step 2: Compute the residual $r_p^{l-1} = f_p^{l-1} - L_p^{l-1} u_p^{l-1}$

step 3: Restrict the residual $r_p^l = (I_{l-1}^l) r_p^{l-1}$

step 4: Restrict the current approximation $v_p^l = (\hat{I}_{l-1}^l) u_p^{l-1}$

for $p = 1$ to P

update 2: Update $v_{p,\delta}^l$ and $r_{p,\delta}^l$ in the external margin cells by using v_p^l and r_p^l of neighboring blocks

for $p = 1$ to P

step 5: Compute the right-hand side $f_{p,\delta}^l = r_{p,\delta}^l + L_{p,\delta}^l v_{p,\delta}^l$

If ($l < M$) then

Apply γ_1 sweeps of

for $p = 1$ to P

step 6: Apply the smoothing method to the system $L_{p,\delta}^l u_{p,\delta}^l = f_{p,\delta}^l$

update 3: Update the overlap cells of neighboring blocks with $u_{p,\delta}^l$

else

step 7: Solve the coarsest-grid problem $L^l u^l = f^l$ on the whole domain

for $p = 1$ to P

update 4: Update $u_{p,\delta}^l$ in the external margin cells by using u_p^l of neighboring blocks

PROLONGATION PART

for $l = M - 1$ to 1

for $p = 1$ to P

step 8: Correct the current approximation $u_p^l = u_p^l - (I_{l+1}^l)(u_p^{l+1} - v_p^{l+1})$

for $p = 1$ to P

update 5: Update $u_{p,\delta}^l$ in the external margin cells by using u_p^l of neighboring blocks.

Apply γ_2 sweeps of

for $p = 1$ to P

step 9: Apply the smoothing method to the system $L_{p,\delta}^l u_{p,\delta}^l = f_{p,\delta}^l$

update 6: Update the overlap cells of neighboring blocks using $u_{p,\delta}^l$.

At step 7, the problem is solved exactly on the coarsest grid common to all the blocks. The cost of solving the system on this grid is negligible compared with the total work.

At updates 2, 4, and 5, the external margin of overlapping cells of neighboring blocks are updated by using inner cells of the current block (Figure 8a). At updates 1, 3, and 6, the overlapping cells of neighboring blocks are updated by using inner and external overlapping cells of the current block (Figure 8b).

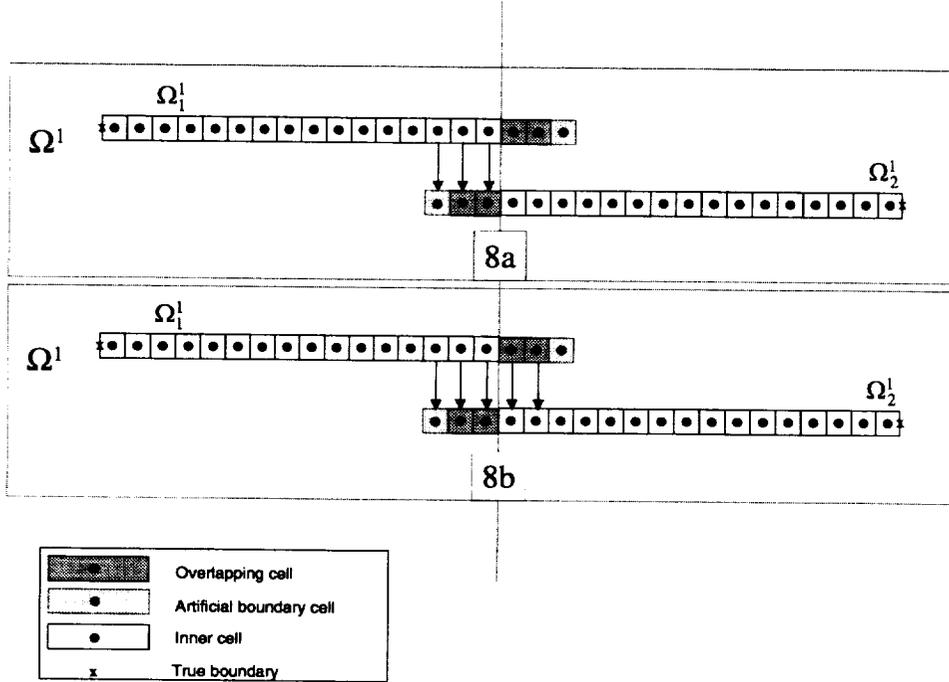


FIG. 8. The algorithm presents two different update processes.

The algorithm includes two computation-update processes:

- Global computation of the residual, restriction, and prolongation (steps 2, 3, 4, and 8) are performed on inner cells of Ω_p^l for all the blocks. Then the external margin of overlap cells of all the blocks are simultaneously updated (updates 2, 4, and 5).
- Within a smoothing step (steps 1, 6, and 9), the update of the solution values at the internal and external overlap cells of neighboring blocks (updates 1, 3, and 6) is performed immediately after completing the smoothing process at the current block $\Omega_{p,\delta}^l$. The new, smoothed $u_{p,\delta}^l$ approximation is used. Thus, the order of these updates repeats the order in which the blocks are treated in the smoothing steps.

Next, we study the properties of the plane-smoother technique for multiblock grids in two situations:

- Anisotropic equation discretized on uniform grids
- Isotropic equation discretized on stretched grids

3. Anisotropic equation discretized on uniform grids. A cell-centered discretization of equation (2.1) on a uniform grid is given by

$$(3.1) \quad \begin{aligned} Lu_{i_x, i_y, i_z} &\equiv \frac{\epsilon_1}{h_x^2} (u_{i_x-1, i_y, i_z} - 2u_{i_x, i_y, i_z} + u_{i_x+1, i_y, i_z}) + \\ &\quad \frac{\epsilon_2}{h_y^2} (u_{i_x, i_y-1, i_z} - 2u_{i_x, i_y, i_z} + u_{i_x, i_y+1, i_z}) + \\ &\quad \frac{1}{h_z^2} (u_{i_x, i_y, i_z-1} - 2u_{i_x, i_y, i_z} + u_{i_x, i_y, i_z+1}) = f'_{i_x, i_y, i_z} \end{aligned}$$

where $i_x = 1, \dots, n_x, i_y = 1, \dots, n_y, i_z = 1, \dots, n_z$, $\epsilon_1 = a/c$, and $\epsilon_2 = b/c$ are the anisotropy coefficients; f' is a discretization of the continuous function $f(x, y, z)/c$; and h_x, h_y , and h_z are the mesh sizes in the x, y , and z directions, respectively. The equation is cell centered and the grid nodes are at the cell vertices. The grid may have different mesh sizes in each dimension.

A multigrid method separates the treatment of different error components: the high-frequency components of the error are reduced in relaxation (smoothing) steps and the low-frequency error components are eliminated in the coarse-grid correction. Thus, the efficiency of a multigrid solver depends strongly on the smoothing factor of the relaxation. Indeed, the derivation of a good smoother is probably the most important step when developing multigrid algorithms.

Consequently, the efficiency of our multigrid technique is determined by the smoothing properties of the blockwise plane smoother. Of course, the smoothing factor of the blockwise plane smoother approaches the factor of a pointwise smoother when the number of blocks increases. In the limit, each block consists of one cell only and the blockwise smoother reverts to a pointwise smoother. This is not a realistic case. We always assume that the number of cells per block is large. Our goal is to study the behavior of the blockwise smoothers when the number of subdomains crossed by the line of a strong anisotropy is low (up to four). This is common in CFD simulations. It is unlikely that regions of very strong anisotropies will traverse many blocks, especially for stretched grids. In such cases, textbook convergence rates could be obtained by allowing larger overlaps in the blocks with the strong anisotropy, while in other blocks the overlaps can be small. Even when an extremely strong anisotropy traverses four blocks in a line, we obtain good convergence rates when using a moderate overlap.

The smoothing properties of a blockwise plane smoother are analyzed for the model equation (3.1) for an anisotropy (in the equation coefficients) ranging from 1 to 10^6 . The number of subdomains crossed by the strong anisotropy is either two or four. As a point of reference, we take the convergence rate obtained by a V(1,0)-cycle for the isotropic problem on a single-block grid with a Gauss-Seidel plane smoother. This cycle demonstrates a textbook multigrid convergence rate of about 0.36 (the plane-smoother convergence of reference) [11]. The convergence rate of the same V(1,0)-cycle with the pointwise Gauss-Seidel smoother is about 0.55 (the point-smoother convergence of reference).

3.1. Numerical results. Let us assume an N^3 global grid partitioned in m^3 blocks (cubic partitioning). Each block is covered with $\Omega_{p,\delta}^l$ extended subgrid. Our aim is to obtain the dependence of the convergence rate $\rho(m, n, \delta, \epsilon)$ on the following parameters:

- number of blocks per side: m
- number of cells per block side: n ($n = \frac{N}{m}$)
- overlap parameter describing intersection between neighboring blocks: δ
- strength of the anisotropy: ϵ

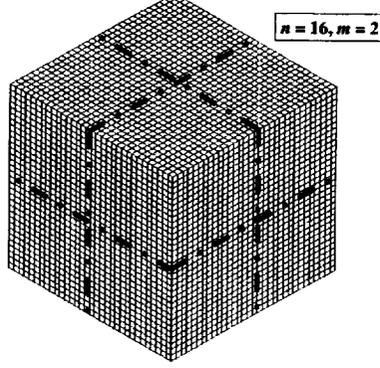


FIG. 9. Rectangular uniform multiblock grid used in the numerical experiments; n is the number of cells per side in every block and m is the number of blocks per side in the split.

3.1.1. Description of test problem. All the numerical experiments reported here deal with the numerical solution of equation (2.1) on the unit cube $\Omega = (0, 1) \times (0, 1) \times (0, 1)$ with the right-hand side of

$$(3.2) \quad f(x, y, z) = -(a + b + c) \sin(x + y + z)$$

discretized on 64^3 and 128^3 multiblock grids (Figure 9) with different overlaps ($\delta = 0, 2, 4,$ and 8). Dirichlet boundary conditions are enforced on the boundary by evaluating

$$(3.3) \quad u(x, y, z) = \sin(x + y + z).$$

The asymptotic convergence rate monitored in the numerical tests is the asymptotic rate of reduction of the L_2 norm of the residual function in one V(1,1)-cycle.

The (x, y) -plane Gauss-Seidel smoother is applied inside each block and the blocks are updated in lexicographic order. For intergrid transfers, the same operators as defined in section 2.1 are used. The 2-D problems defined in each plane are solved by one 2-D V(1,1)-cycle with x -line Gauss-Seidel smoothing.

3.1.2. Discussion of results. Two sets of calculations were performed to determine the experimental asymptotic convergence rate as a function of the anisotropies. In the first set, two parameters (ϵ_1 and ϵ_2) are varied, and in the second, just one parameter (ϵ_1) is varied while the second parameter (ϵ_2) remains equal to 1.

- Both ϵ_1 and ϵ_2 are varied for the single block case ($m = 1$) and two multiblock cases ($m = 2$ and $m = 4$). Different overlaps between blocks ($\delta = 0, 1,$ and 2) are examined. The upper graphs in Figures 10 and 11 show the results for 64^3 ($N = 64$) and 128^3 ($N = 128$) grids, respectively.
- Only the parameter ϵ_1 is varied for the single block case ($m = 1$), and two multiblock cases ($m = 2$ and $m = 4$). Different overlaps between blocks ($\delta = 0, 1,$ and 2) are examined. The lower graphs in Figures 10 and 11 show these results for 64^3 ($N = 64$) and 128^3 ($N = 128$) grids respectively.

All the graphs exhibit a similar behavior with respect to δ and ϵ . We can distinguish three different cases:

- In the single-block case ($m = 1$), the convergence rate decreases quickly for an anisotropy larger than 100, tending to (nearly) zero for very strong anisotropies. In fact, the convergence rate per V(1,1)-cycle decreases quadratically with increasing anisotropy strength, as was predicted in earlier work [11].

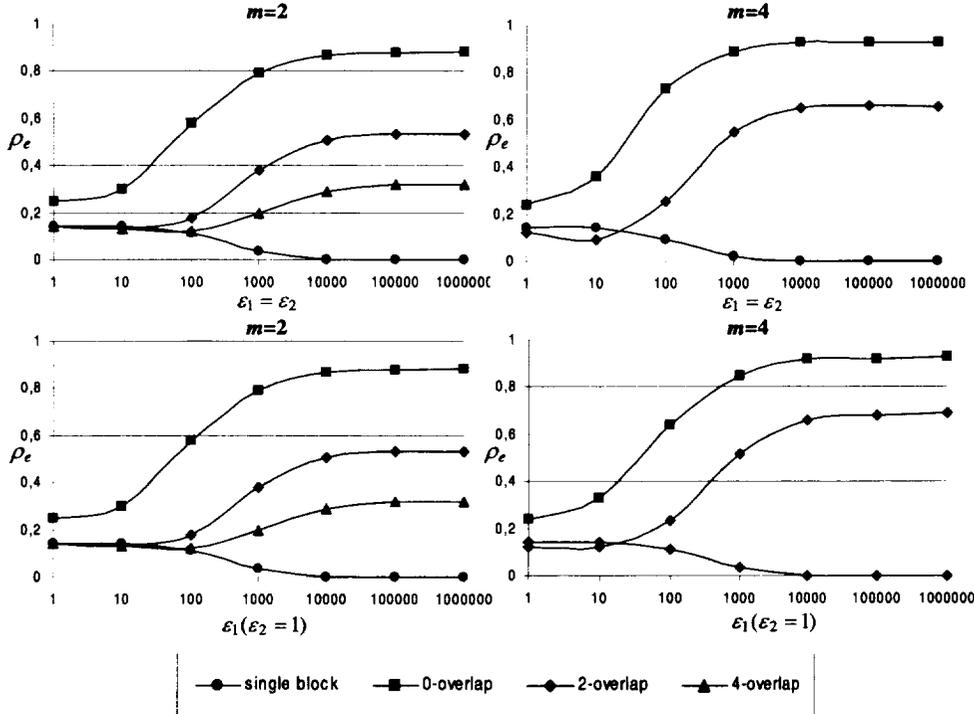


FIG. 10. Experimental convergence factors, ρ_e , of one 3-D $V(1,1)$ -cycle with blockwise (x, y) -plane smoother in a 64^3 grid with respect to the anisotropy strength (ϵ_1 and ϵ_2) for different values of the overlap (δ) and the number of blocks per side (m).

- If the domain is blocked with the minimal overlapping ($m > 1$ and $\delta = 0$), the convergence rate for small anisotropies ($1 \leq \epsilon \leq 100$) is similar to that obtained with a single block with a pointwise smoother on the whole domain (i.e., point-smoother convergence of reference, which is about 0.55^2). It increases to a fixed value for larger anisotropies. This limit is defined by the convergence in the corresponding domain decomposition algorithm. The convergence rate for strong anisotropies gets closer to one for finer grids (compare Figures 10 and 11 for the same m and δ but different N).
- If the domain is blocked with larger overlapping ($m > 1$ and $\delta \geq 2$), the convergence rate for small anisotropies is similar to that obtained without blocking on the whole domain (i.e. plane-smoother convergence of reference which is about 0.36^2) and increases to the fixed domain decomposition limit for very strong anisotropies. The asymptotic value for strong anisotropies gets closer to one for smaller overlaps and finer grids (compare Figures 10 and 11 for the same m but different δ and N).

Numerical results show the following properties of the convergence behavior:

- Convergence is very poor with the minimal overlap ($\delta = 0$), but it improves rapidly with larger overlaps ($\delta \geq 2$).
- If we compare the results for 64^3 and 128^3 and $m = 2$ (Figures 10 and 11), we realize that the convergence rate for large anisotropies is bounded if the overlap is increased in proportion to the size of the block, i.e., $\rho(2, n, \delta, \epsilon) \approx \rho(2, 2n, 2\delta, \epsilon)$ (see Figure 12). Therefore, mesh-independent convergence rates can be obtained for large anisotropies by increasing the number of overlap cells in proportion to the number of cells per side in the subgrid. The amount of extra work (in relative

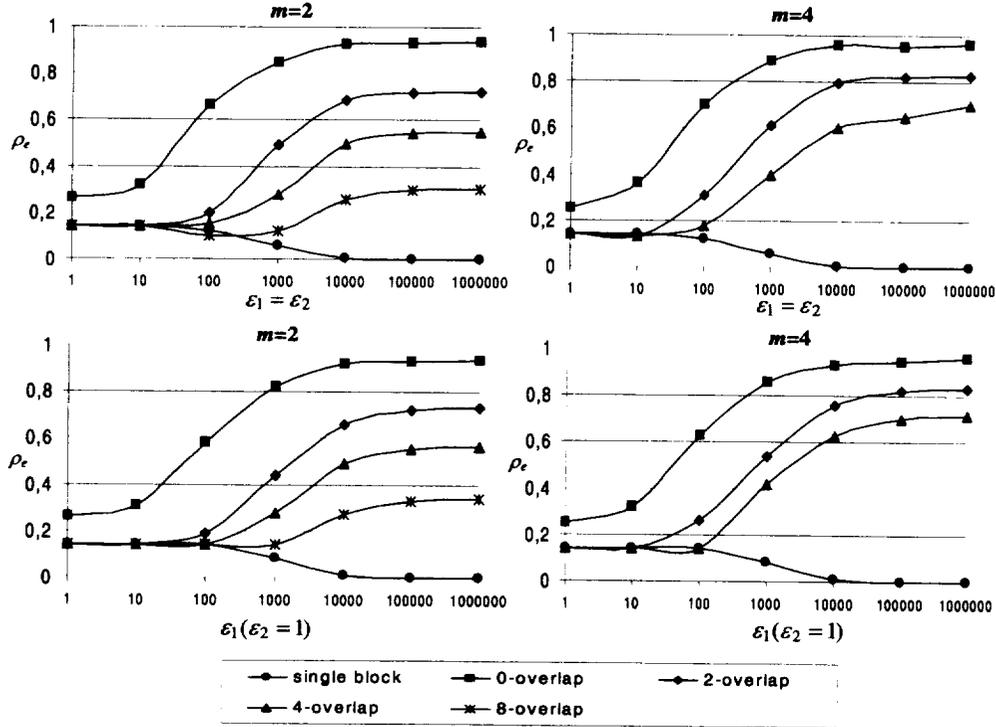


FIG. 11. Experimental convergence factors, ρ_e , of one 3-D $V(1,1)$ -cycle with blockwise (x, y) -plane smoother in a 128^3 grid with respect to the anisotropy strength (ϵ_1 and ϵ_2) for different values of the overlap (δ) and the number of blocks per side (m).

terms) to maintain the convergence rate at a fixed level does not increase for finer grids.

- For 17-cell overlap ($\delta = 8$), 2 blocks per side ($m = 2$), and 32^3 cells per block ($n = 32$), we obtain for all anisotropies (nearly) the same convergence rate as for the isotropic case without blocking (plane-smoother convergence of reference). This overlap results in only a 25 percent computation and memory penalty.
- The smoothing factor for very strong anisotropies degenerates with an increasing number of blocks m , because the blockwise smoother tends to be a pointwise smoother.

3.2. Analysis. The full space Fourier analysis is known as a simple and very efficient tool to predict convergence rates of multigrid cycles in elliptic problems. However, this analysis is inherently incapable of accounting for boundary conditions. In isotropic elliptic problems where boundary conditions affect just a small neighborhood near the boundary, this shortcoming does not seem to be very serious. The convergence rate of a multigrid cycle is mostly defined by the relaxation smoothing factor in the *interior* of the domain, and therefore, predictions of the Fourier analysis prove to be in amazing agreement with the results of numerical calculations. In the presence of a strong anisotropy, boundary conditions have a stronger effect on the overall convergence factor that a multigrid cycle can achieve. If the strong anisotropy direction aligns with the grid (as in equation (3.1)) and the domain is not blocked, then the smoothing factor calculated by the full-space Fourier analysis still predicts well the multigrid cycle convergence rate. In the case of nonalignment, however, the full-space analysis cannot reflect the poorness of the coarse-grid correction, which slows down the convergence (see [5, 3]) and so the half-space analysis must be used instead. The

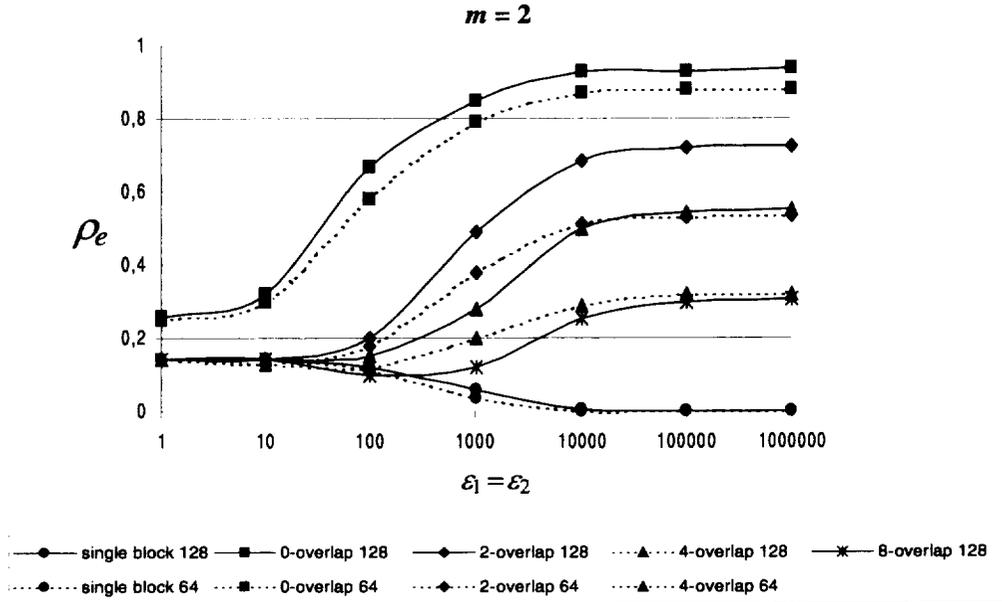


FIG. 12. *Experimental convergence factors, ρ_e , of one 3-D $V(1,1)$ -cycle with blockwise (x, y) -plane smoother in 128^3 and 64^3 grids with respect to the anisotropy strength (ϵ_1 and ϵ_2) for different values of the overlap (δ) and 2 blocks per side ($m = 2$).*

multiblock structure brings additional difficulties: in strongly anisotropic problems on decomposed domains, artificial boundaries have a large effect on the solution in the interior (e.g., on the relaxation stage), making even the half-space analysis inadequate. In the extreme case of a very large anisotropy, the behavior of a plane smoother is similar to that exhibited by a one-level overlapping Schwartz method [1]. The following analysis is an extension of the half-space Fourier analysis, taking into account the influence of a second boundary.

We consider the discrete equation (3.1) on a layer $(i_x, i_y, i_z) : 0 \leq i_x \leq N, -\infty < i_y, i_z < \infty$. This domain is decomposed by overlapped subdomains in an x -line partitioning. The boundaries of all the subdomains are given by the set of cell planes orthogonal to the x -coordinate axis. Boundary conditions are represented by one 2-D Fourier component at a time. In this way, the original 3-D problem is translated into a 1-D problem where frequencies of the Fourier component are considered parameters. When estimating the smoothing factor for the 3-D alternating-direction plane smoother, we analyze only the high-frequency Fourier components. A 2-D Fourier component $e^{i(\theta_y i_y + \theta_z i_z)}$, ($|\theta_y| \leq \pi; |\theta_z| \leq \pi$) is a high-frequency component if the absolute value of either θ_y or θ_z is greater than or equal to $\pi/2$.

We are given values of $N, m, n, \delta, \epsilon_1$ and ϵ_2 (see definitions in section 3.1) and we let the integers $i_x = 0 \dots N$ numerate the cells. We also label the blocks with numbers from 1 to m . To predict the smoothing factors observed in our numerical experiments, we assume that the strongest anisotropy direction is aligned with the x axis, i.e., $\epsilon_1 > \epsilon_2 \geq 1$. The actual boundary separating two neighboring blocks is placed *between* the cell centers. Recall that the overlap parameter δ corresponds to the case where the width of the domain

shared by two overlapped blocks is $(2\delta + 1)h$. The left boundary condition of the first block is defined at the center of the actual (not artificial) left boundary cell $l_1 = 0$; the right boundary condition is defined at $r_1 = n + \delta + 1$. Each block k ($k = 2, \dots, m - 1$) has two (artificial) boundary planes settled in the interior of the domain: the center of the left boundary cell is located at $l_k = (k - 1)n - \delta$, and the center of the right boundary cell is at $r_k = kn + \delta + 1$. The left boundary cell of the last m -th block is centered at $l_m = (m - 1)n - \delta$, and the center of the right boundary cell is at $r_m = N$. One application of the 3-D alternating-direction plane smoother on the k -th block includes three plane-relaxation sweeps. In the analysis, we assume that in the first sweep, the smoother proceeds in the x direction from $i_x = l_k$ to $i_x = r_k$, solving the corresponding y - z -plane problems exactly. In numerical tests, just one V(1,1)-cycle is applied to approximate the 2-D solutions. The second and third sweeps are performed in the y and z directions respectively.

Let us consider equation (3.1) with the initial approximation $\tilde{u}_{i_x, i_y, i_z}$ and the source function f'_{i_x, i_y, i_z} given in the form

$$\begin{aligned}\tilde{u}_{i_x, i_y, i_z} &= A(i_x)e^{i(\theta_y i_y + \theta_z i_z)} \\ f'_{i_x, i_y, i_z} &= F(i_x)e^{i(\theta_y i_y + \theta_z i_z)} \\ i_x &= 0, \dots, N,\end{aligned}$$

where A and F are $(N + 1)$ dimensional complex-valued vectors. $A(i_x)$ and $F(i_x)$ represent the corresponding amplitudes of the Fourier component at i_x .

Let $E(i_x)$ be the amplitude of the algebraic error in the k -th block after updating the solution values at the overlap $[l_k, (k - 1)n]$. We have $E(i_x) = U(i_x) - A(i_x)$, where $U(i_x)$ is the amplitude of the exact discrete solution

$$\begin{aligned}U(0) &= U_0; & U(N) &= U_1, \\ \epsilon_1 U(i_x - 1) + \mu_1 U(i_x) + \epsilon_1 U(i_x + 1) &= F(i_x), \\ i_x &= 1, \dots, N - 1, \\ \mu_1 &= -2\epsilon_1 + 2\epsilon_2(\cos \theta_y - 1) + 2(\cos \theta_z - 1)\end{aligned}$$

with given boundary conditions U_0 and U_1 .

In the first sweep, the new amplitude $E_1(i_x)$ is defined from

$$\begin{aligned}(3.4) \quad E_1(l_k) &= E(l_k); & E_1(r_k) &= E(r_k), \\ \epsilon_1 E_1(i_x - 1) + \mu_1 E_1(i_x) &= -\epsilon_1 E(i_x + 1), \\ i_x &= l_k + 1, \dots, r_k - 1.\end{aligned}$$

After the second sweep the new amplitude $E_2(i_x)$ is the solution of the following system of equations.

$$\begin{aligned}(3.5) \quad E_2(l_k) &= E(l_k), & E_2(r_k) &= E(r_k), \\ \epsilon_1 E_2(i_x - 1) + \mu_2 E_2(i_x) + \epsilon_1 E_2(i_x + 1) &= -\epsilon_2 e^{i\theta_y} E_1(i_x), \\ i_x &= l_k + 1, \dots, r_k - 1, \\ \mu_2 &= -2\epsilon_1 + \epsilon_2(e^{-i\theta_y} - 2) + 2(\cos \theta_z - 1)\end{aligned}$$

Finally, the amplitude $E_3(i_x)$ of the approximate solution obtained after the third sweep is found from the following system of equations.

$$E_3(l_k) = E(l_k), \quad E_3(r_k) = E(r_k),$$

$$\begin{aligned}
(3.6) \quad & \epsilon_1 E_3(i_x - 1) + \mu_3 E_3(i_x) + \epsilon_1 E_3(i_x + 1) = -e^{i\theta_x} E_2(i_x), \\
& i_x = l_k + 1, \dots, r_k - 1, \\
& \mu_3 = -2\epsilon_1 + \epsilon_2 (\cos \theta_y - 1) + (e^{-i\theta_x} - 2).
\end{aligned}$$

The smoothing factor of the alternating-direction plane smoother step is the ratio between the L_2 norms of high-frequency errors before and after the step is performed. This analysis is very rigorous and especially useful for predicting the convergence in full multigrid (FMG) algorithms (see, e.g., [2]). In these algorithms, where only a few smoothing steps are performed on each grid and the initial error is basically the difference between solutions on successive grids, this analysis allows a direct estimation of the algebraic error at any stage of the algorithm.

To estimate asymptotic convergence rates, the following simplifications can be adopted.

3.2.1. Simplifications in different regimes. There are two processes affecting the amplitude of high-frequency error components. The first is the smoothing in the interior of the blocks. The effect of this smoothing is very similar to that on the single-block domain. It is well described by the smoothing factor SM derived from the usual full-space Fourier mode analysis (see [2, 11, 20]). If the problem is essentially isotropic ($\epsilon_1 = O(1)$), then each of the three sweeps (3.4)–(3.6) contributes to the alternating-direction plane-relaxation smoothing factor. This smoothing factor proves to be very small ($SM = 1/\sqrt{125} \approx 0.089$ in the pure isotropic problem), but the overall convergence rate in a V-cycle is worse because it is dictated by the coarse-grid correction. To predict the asymptotic convergence rate in a V-cycle, the two-level analysis should be performed. In problems with a moderate anisotropy in the x direction ($\epsilon_1 = O(h^{-1})$), the high-frequency error is reduced mainly in the third sweep (3.6) and the smoothing factor approaches the 1-D factor $SM = 1/\sqrt{5}$ (see [11]). In strongly anisotropic problems ($\epsilon_1 = O(h^{-2})$), the third sweep reduces the smooth error components as well, actually solving the problem rather than just smoothing the error.

The second process influencing the high-frequency error is the error propagation from incorrectly specified values at the block boundaries. The distance on which this high-frequency boundary error penetrates inside blocks strongly depends on the anisotropy. It can be estimated by considering a semi-infinite problem associated with the sweep (3.6). The left-infinite problem stated for the k -th block assumes a zero right-hand side and omits the left boundary condition

$$E(r_k) = B, \quad E_2(i_x) = 0, \quad -\infty < i_x < r_k.$$

The solution of this problem is

$$E_3(i_x) = B \lambda_l(\theta_y, \theta_z)^{i_x - r_k},$$

where $\lambda_l(\theta_y, \theta_z)$ satisfies

$$(3.7) \quad \epsilon_1 \lambda^{-1} + \mu_3 + \epsilon_1 \lambda = 0; \quad |\lambda| \geq 1.$$

For high-frequencies under the assumption $\epsilon_1 > \epsilon_2 \geq 1$,

$$\left| \lambda_l(\theta_y, \theta_z) \right| \geq \Lambda_l = \left| \lambda_l(0, \pi/2) \right|.$$

The right-infinite problem is similarly solved providing

$$\left| \lambda_r(\theta_y, \theta_z) \right| \leq \Lambda_r = \left| \lambda_r(0, \pi/2) \right|,$$

where $\lambda_r(\theta_y, \theta_z)$ is a root of equation (3.7) satisfying $|\lambda| \leq 1$.

The roots of quadratic equation (3.7) for $\theta_y = 0$ and $\theta_z = \frac{\pi}{2}$ are given by

$$\Lambda_l = \left| -\frac{\hat{\mu}_3}{2} + \sqrt{\frac{\hat{\mu}_3^2}{4} - 1} \right| \quad \text{and}$$

$$\Lambda_r = \Lambda_l^{-1} = \left| -\frac{\hat{\mu}_3}{2} - \sqrt{\frac{\hat{\mu}_3^2}{4} - 1} \right|,$$

where

$$\hat{\mu}_3 = -2 - \frac{(i+2)}{\epsilon_1}.$$

Thus, if

$$\Lambda_r^{r_k - l_k - 2\delta - 1} \ll \Lambda_r^{2\delta + 1},$$

i.e., the influence of the far boundary (e.g., r_{k+1} boundary on r_k), is negligible in comparison with the influence of the nearby boundary (l_{k+1} on r_k), then the high-frequency error amplitude reduction factor, RF , is estimated as the ratio between the amplitudes on the right boundary r_k before and after the plane-smoother step is performed. After relaxing the k -th block, the amplitude at l_{k+1} is about $B\Lambda_l^{-2\delta-1}$ and after relaxing the $(k+1)$ -th block, the estimated amplitude at r_k is $B(\Lambda_r/\Lambda_l)^{2\delta+1} = B(\Lambda_r)^{4\delta+2}$. Thus, the reduction factor could be approximated by

$$(3.8) \quad RF \approx (\Lambda_r)^{4\delta+2}.$$

If the anisotropy is strong ($\epsilon_1 = O(h^{-2})$), both boundaries (far and nearby) affect the error amplitude reduction factor. If the number of blocks m is not too large, then the corresponding problem includes m coupled problems (like 3.6) with zero right-hand sides. This multiblock problem can directly be solved. For the two-block partition it results in

$$(3.9) \quad RF = \left(\frac{\Lambda_l^{n-\delta-1} - \Lambda_r^{n-\delta-1}}{\Lambda_l^{n+\delta} - \Lambda_r^{n+\delta}} \right)^2.$$

In all the numerical tests described in sections 3.1–3.2, the V(1,1)-cycle convergence rate can be predicted by the (squared) smoothing factor, which is calculated as the maximum value of SM or RF .

3.2.2. Relation to domain decomposition convergence theory. When the anisotropy is very strong and the number of blocks is much greater than two, convergence rates observed in numerical tests are well described by the domain decomposition convergence theory [1]. In this case, we are not interested in the smoothing properties of the 3-D smoother but in its resolution properties, because it becomes an exact method to solve N noncoupled 2-D problems (when ϵ_1 and ϵ_2 are very large) or N^2 1-D problems (when only ϵ_1 is large). In fact, asymptotic convergence rates for large anisotropies are in good agreement with convergence rates predicted for 1-D or 2-D domain-decomposition Schwartz methods.

3.2.3. Comparison with numerical tests. For simplicity, we consider the V(1,1)-cycle with only z -planes used in the smoothing step. The assumption that $\epsilon_1 > \epsilon_2 \geq 1$ validates this simplification. The predicted smoothing factor for this relaxation is $SM = 1/\sqrt{5}$ [11]. In numerical calculations for isotropic problems on single-block domains, the asymptotic convergence rate was 0.14 per V(1,1)-cycle, which is very close to the value predicted by the two-level Fourier analysis (0.134).

TABLE 3.1

Experimental, ρ_e , and analytical, ρ_a , convergence factors of a single 3-D $V(1,1)$ -cycle with blockwise (x,y) -plane Gauss-Seidel smoother ($2 \times 2 \times 2$ partition) versus anisotropy strength, width of the overlap and the block size

| ϵ_1 | δ | $n = 64$ | | $n = 128$ | |
|--------------|----------|----------|----------|-----------|----------|
| | | ρ_e | ρ_a | ρ_e | ρ_a |
| 10^6 | 0 | 0.882 | 0.881 | 0.939 | 0.939 |
| | 2 | 0.534 | 0.529 | 0.731 | 0.729 |
| | 4 | 0.321 | 0.316 | 0.567 | 0.566 |
| | 8 | | | 0.341 | 0.340 |
| 10^4 | 0 | 0.87 | 0.87 | 0.92 | 0.92 |
| | 2 | 0.51 | 0.51 | 0.66 | 0.67 |
| | 4 | 0.29 | 0.29 | 0.49 | 0.49 |
| | 8 | | | 0.27 | 0.26 |
| 10^2 | 0 | 0.58 | 0.56 | 0.58 | 0.56 |
| | 2 | 0.18 | 0.14 | 0.19 | 0.14 |
| | 4 | 0.12 | 0.14 | 0.14 | 0.14 |
| | 8 | | | 0.14 | 0.14 |

The next case we present is the comparison of the analytical predictions and the numerical results for the $V(1,1)$ -cycle convergence rates for different anisotropies for the domain decomposed into two blocks in each direction. In this case, for predicting the reduction factor RF , analytical expression (3.9) can be used. The formula for the asymptotic convergence rate ρ_a is

$$(3.10) \quad \rho_a = \max(RF^2, 0.14).$$

Table 3.1 presents a representative sample of experiments for the case $\epsilon_1 \gg \epsilon_2 = 1$. In this table, ρ_e corresponds to asymptotic convergence rates observed in the numerical experiments, while ρ_a is calculated by means of equation (3.10). The results demonstrate nearly perfect agreement.

The asymptotic convergence rate deteriorates as the number of blocks grows. For a multiblock partition, an accurate value of RF can still be calculated. (In the case when the domain is partitioned into m blocks, it will be the spectral radius of a $(m-1)$ -by- $(m-1)$ matrix.) We decided, however, to present another, simplified, methodology. Here the convergence upper bound, ρ_{DD} , which is the asymptotic convergence rate in the multiblock domain decomposition solver is estimated numerically as $\rho_{DD} = \rho_e$ for $\epsilon_1 = \infty$. The reduction factor RF is calculated from expression (3.8), and the predicted convergence rate is

$$(3.11) \quad \rho_a = \max(\min(RF^2, \rho_{DD}), 0.14).$$

Practically, this formula implies that for small anisotropies ($\epsilon_1 \leq 10$) and for very large anisotropies ($\epsilon_1 \geq 10^4$), the asymptotic convergence rate is estimated by some constants calculated numerically. The analytical expression is correct in the range of moderate anisotropies. Notice that RF depends only on two parameters (ϵ_1 and δ) which means that estimate (3.11) provides a good prediction only if $n \gg \delta$. Table 3.2 compares this analytical prediction with the results of numerical tests for some moderate anisotropies and different values of the overlap parameter δ .

TABLE 3.2

Rectangular stretched multiblock grid used in the numerical experiments; n is the number of cells per side in every block, m is the number of blocks per side in the split, and α is the stretching ratio.

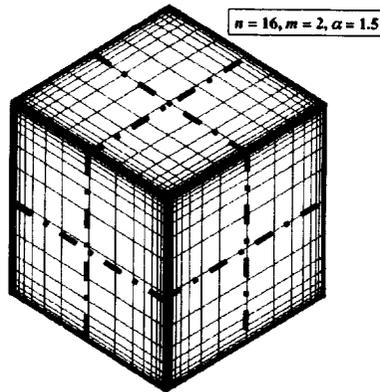
| ϵ_1 | δ | ρ_a | ρ_e | | | | |
|--------------|----------|----------|----------|---------|----------|---------|----------|
| | | | $n = 16$ | | $n = 32$ | | $n = 64$ |
| | | | $m = 4$ | $m = 2$ | $m = 4$ | $m = 2$ | |
| 10^2 | 0 | 0.56 | 0.64 | 0.58 | 0.63 | 0.58 | |
| | 2 | 0.14 | 0.23 | 0.18 | 0.26 | 0.19 | |
| 10^3 | 0 | 0.83 | 0.85 | 0.79 | 0.86 | 0.82 | |
| | 2 | 0.40 | 0.52 | 0.38 | 0.54 | 0.44 | |
| | 4 | 0.19 | | 0.20 | 0.42 | 0.28 | |

TABLE 3.3

Moderate anisotropy: experimental, ρ_e , and analytical, ρ_a , convergence factors versus overlap parameter, block size, and number of blocks

We can now discuss the different convergence behavior exhibited by fixed and decreasing overlaps. A decreasing overlap introduces errors on the coarse grids which become smooth errors on the fine grids and, therefore, cannot be eliminated by the smoother on the fine grids. This is especially true in cases of weak to moderate anisotropy when the reduction factor is defined by smoothing properties rather than by the properties of a domain decomposition solver. The convergence in this case is still bounded above by the convergence rate of the domain decomposition solver, but we lose a lot of efficiency, especially because the domain decomposition solver is sensitive to such issues as number of blocks and overlap size.

For a given m and δ/n ratio, if the domain decomposition solver is efficient enough, then it is possible to use the same δ/n ratio on coarser grids rather than using a fixed number of cells in the overlap regions. However, adopting a fixed δ/n ratio as a general rule for overlapping reduces the solver down to the level of a domain decomposition solver. In the case of fixed overlap, the domain decomposition solver efficiency is just an upper bound on the convergence rate. The bound is sharp only for huge anisotropies (huge in strength and relevant region size).



4. Isotropic equation discretized on stretched grids. In this section, we study numerically the behavior of plane-implicit smoothers for multiblock grids when the anisotropy is not uniform but exponential. For stretched grids, each cell can have a unique level of anisotropy (see Figure 13). The stretching is applied near the boundary to mimic grids used in CFD simulations.

Consider an N^3 global grid partitioned into m^3 blocks (cubic partitioning)(Figure 13). Again, each block is supplied with an extended $\Omega_{p,\delta}^l$ subgrid. Our goal is to analyze the dependence of the convergence rate $\rho(m, n, \delta, \alpha)$ on the following parameters:

- number of blocks per side: m
- number of grid cells per block side: n
- overlap parameter: δ
- stretching ratio: $\alpha = \frac{h_{i+1}}{h_i}$

4.1. Description of test problem. All the numerical experiments reported here deal with the numerical solution of equation (2.1) on the unit cube $\Omega = (0, 1) \times (0, 1) \times (0, 1)$ with a right-hand side and boundary conditions given by expressions (3.2) and (3.3). The discretization is done on 64^3 and 128^3 rectangular stretched multiblock grids with various overlaps ($\delta = 0, 2, 4, 8$).

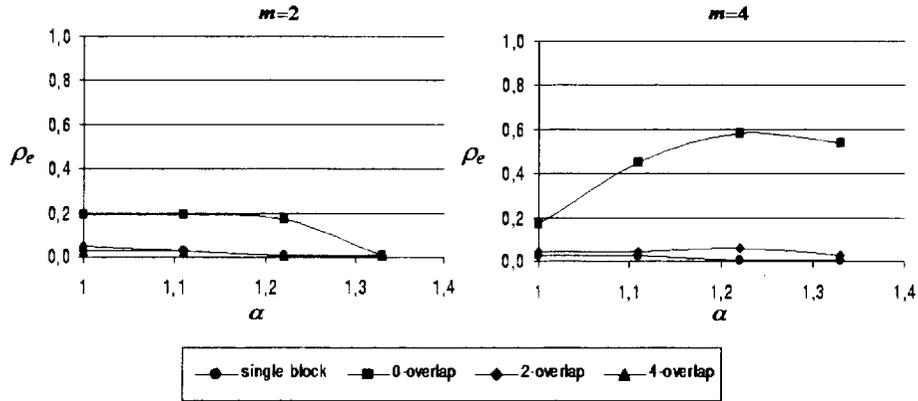


FIG. 13. Experimental convergence factors, ρ_e , of one 3-D $V(1,1)$ -cycle with blockwise alternating-direction plane smoother in a 64^3 grid with respect to the stretching ratio (α) for different values of the overlap (δ) and number of blocks per side (m).

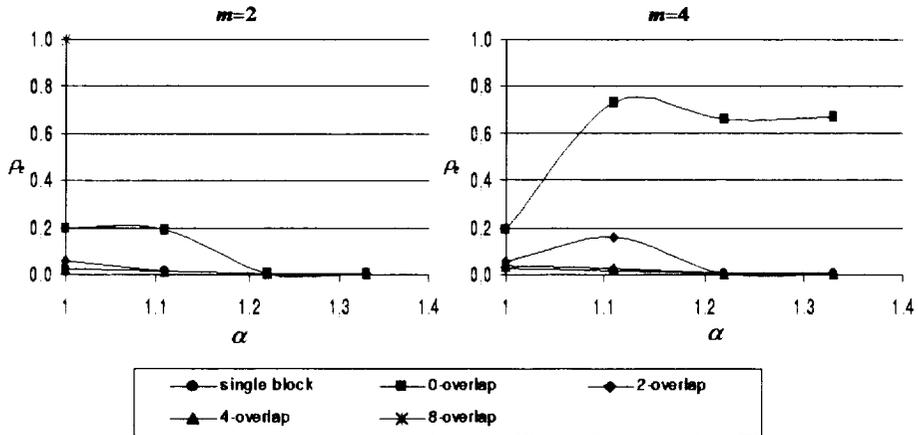


FIG. 14. Experimental convergence factors, ρ_e , of one 3-D $V(1,1)$ -cycle with blockwise alternating-direction plane smoother in a 128^3 grid with respect to the stretching ratio (α) for different values of the overlap (δ) and number of blocks per side (m).

As before, we tested a FAS version of $V(1,1)$ -cycle with the blockwise alternating-direction plane smoother, unweighted solution averaging, volume-weighted residual averaging, and trilinear correction interpolation. The solutions of the 2-D problems in each plane are approximated by one 2-D $V(1,1)$ -cycle

employing alternating-direction line Gauss-Seidel smoothing. The asymptotic convergence rates observed in the numerical tests are presented in Figures 14 and 15.

4.2. Discussion of results. Numerical simulations were performed to obtain the experimental convergence rate with respect to the stretching ratio, α . The single-block and multiblock grids ($m = 1, 2$, and 4) with different overlaps ($\delta = 0, 2$, and 4) were tested. Figure 14 shows the results for a 64^3 grid, and Figure 15 for a 128^3 grid. The results can be summarized in the following two observations:

- With a 2^3 partitioning, even the minimum overlap ($\delta = 0$) is enough to exhibit good convergence rates. The results for a multiblock grid with overlap of $\delta = 2$ match with the results obtained for the single-block anisotropic case. That is, the convergence rate tends toward zero as the anisotropy increases.
- With a 4^3 partitioning, results are not as good. With the minimal overlap ($\delta = 0$), the convergence rate degrades for finer grids. However, with a larger overlap ($\delta = 2$), the convergence rate again tends towards the convergence rate demonstrated in single-block grid anisotropic cases.

5. Parallel implementation of multigrid technique. Block-wise plane smoothers may also be used to facilitate the parallel implementation of a multigrid method on a single-block (rectangular) grid. Notice that in this case there are global lines and planes and we are interested in block-wise smoothers only for purposes of parallel computing. To get a parallel implementation of a multigrid method, one can adopt one of the following strategies (see, e.g., [12]).

- *Domain decomposition:* A domain decomposition is applied first. Then, a multigrid method is used to solve problems inside each block.
- *Grid partitioning:* A multigrid method is used to solve the problem in the whole grid.

Domain decomposition is easier to implement and implies fewer communications (better parallel properties), but it has a negative impact on the convergence rate. On the other hand, grid partitioning implies more communication but it retains the convergence rate of the sequential algorithm (better numerical properties).

From a parallel code designer point of view, our approach is somewhere between domain decomposition and grid partitioning. Block-wise plane smoothers appear to be a tradeoff between architectural and numerical properties (domain decomposition and grid partitioning).

- *Convergence rate.* For the isotropic case the convergence rate is equal to that obtained with grid partitioning, and it approaches the convergence rate of a domain decomposition method as the anisotropy becomes stronger.
- *Communications.* Although higher than in domain decomposition, the number of communications is lower than in grid-partitioning algorithms.

However, it should be noted that due to the lack of definition of global planes and lines, grid partitioning is not viable in general multiblock grids.

Therefore, the use of block-wise smoothers is justified to facilitate parallel processing when the problem does not possess a strong anisotropy crossing the whole domain. In such a case, the expected convergence rate (when using moderate overlaps at the block interfaces crossed by a strong anisotropy) is similar to the rate achieved with grid partitioning, but the number of communications is considerably lower.

6. Conclusions and future directions. We have analyzed in detail the smoothing properties of blockwise plane smoothers for multiblock grids. For simplicity, our analysis and test problems have focused on rectangular grids, although the results can be extrapolated to more general multiblock grids. As a point of reference, we take the convergence rate obtained by the plane smoother in the isotropic problem on a

single-block grid.

Typically, in true CFD simulations, there are no very large anisotropies crossing the whole domain. The regions of very strong anisotropies are narrow (usually where there are thin boundary layers). Therefore, it is unlikely that such a region will be divided into many blocks in the direction normal to the thin layer. Anisotropies crossing one, two, or up to four blocks are very practical, especially for stretched grids. For these cases, the overall convergence rate is still very good with a moderate overlap, as has been demonstrated in the present report.

The blockwise plane method works as a smoother for low anisotropies and becomes a domain decomposition solver for very large anisotropies. We have obtained an analytical expression that predicts the smoothing properties of the blockwise plane-implicit smoother with respect to the block size, the amount of overlap, and the strength of the anisotropy. The expression has been validated with numerical experiments.

As an example, for a line of strong anisotropy partitioned into two blocks, the multigrid convergence of reference can be obtained with an overlap of just 25% of the number of cells in the block, and this convergence rate can be maintained by increasing linearly the number of overlapping cells with the problem size. In actual problems, a much smaller overlap is likely to be sufficient.

We have also analyzed the case of an isotropic equation on grids which are stretched near to the boundaries, a common case in practical CFD problems. This case is even more favorable, as the convergence rate obtained for the single-block grid is maintained in multiblock grids with very small overlaps, tending to zero with increasing anisotropy strength.

Blockwise alternating-direction plane relaxation methods are found to be robust smoothers and their use should be considered for inclusion in the next generation of production CFD codes. A robust multiblock code should check whether there is an anisotropy crossing a block interface. If so, an extended subgrid overlapping with the neighboring block should be constructed. A multiblock strategy opens the possibility of using an adaptive smoother — that is, different smoothers for different portions of the domain, where choice of the smoother is based on a minimization of operation count while retaining optimum smoothing performance. An example of this would be using a plane-implicit smoother in the portions of the domain that have strong anisotropies while using a point smoother in the regions that are isotropic.

The present code updates the blocks following the lexicographic ordering. However, in order to run the code on a parallel computer, so that each processor solves a set of blocks, the update ordering must present a higher parallelism grade. The next step will be to analyze the convergence rate and architectural properties of red-black (or general-colored) ordering of the blocks. We intend to continue the work on the block-structured algorithms. In particular, we will study the applicability of blocked alternating-direction plane methods as multigrid smoothers for convection-dominated problems and more complicated PDE's.

7. Acknowledgments. The results of the simulations were obtained on the Silicon Graphics Incorporated Origin 2000 operated by the Aerodynamic and Acoustic Methods Branch at the NASA Langley Research Center.

REFERENCES

- [1] P. B. B. SMITH AND W. GROPP, *Domain Decomposition, Parallel Multilevel Methods for Elliptic Partial Differential Equations*, Cambridge University Press, 1996.
- [2] A. BRANDT, *Multigrid techniques: 1984 guide with applications to fluid dynamics*, Tech. Report GMD-Studien 85, May 1984.

- [3] A. BRANDT AND B. DISKIN, *Multigrid solvers for non-aligned sonic flows*. to appear in SIAM J. Sci. Comp.
- [4] ———, *Multigrid solvers on decomposed domains*, in Domain Decomposition Methods in Science and Engineering, A. Quarteroni, J. Periaux, Y. A. Kuznetsov, and O. Widlund, eds., Contemp. Math., Amer. Math. Soc., 157 (1994), pp. 135–155.
- [5] ———, *Multigrid solvers for the non-aligned sonic flow: The constant coefficient case*, Computers and Fluids, 28 (1999), pp. 511–549. Also Gauss Center Report WI/GC-8 at The Weizmann Institute of Science, Oct. 1997.
- [6] D. J. MAVRIPLIS, *Multigrid strategies for viscous flow solvers on anisotropic unstructured meshes*, in Proceedings of the AIAA CFD Conference, June 1997.
- [7] C. C. DOUGLAS, *Caching in with multigrid algorithms: Problems in two dimensions*, Paral. Alg. and Appl., 9 (1996), pp. 195–204.
- [8] J. E. DENDY JR., M. P. IDA, AND J. M. RUTLEDGE, *A semi-coarsening multigrid algorithm for SIMD machines*, SIAM J. Sci. Stat. Comput., 13 (1992), pp. 1460–1469.
- [9] J. E. JONES AND N. D. MELSON, *A note on multi-block relaxation schemes for multigrid solvers*, ICASE Report No. 97-15, 1997.
- [10] J. LINDEN, G. LONSDALE, H. RITZDORF, AND A. SCHÜLLER, *Scalability aspects of parallel multigrid*, Future Generation Computer Systems, 10 (1994), pp. 429–440.
- [11] I. M. LLORENTE AND N. D. MELSON, *Robust multigrid smoothers for three dimensional elliptic equations with strong anisotropies*, Tech. Report ICASE Report 98-37, 1998.
- [12] I. M. LLORENTE AND F. TIRADO, *Relationships between efficiency and execution time of full multigrid methods on parallel computers*, IEEE Trans. on Parallel and Distributed Systems, 8 (1997), pp. 562–573.
- [13] W. MULDER, *A new multigrid approach to convection problems*, J. Comput. Phys., 83 (1989), pp. 303–323.
- [14] N. H. NAIK AND J. V. ROSENDALE, *The improved robustness of multigrid elliptic solvers based on multiple semicoarsened grids*, SIAM J. of Num. Anal., 30 (1993), pp. 215–229.
- [15] C. W. OOSTERLEE, *A GMRES-based plane smoother in multigrid to solve 3-D anisotropic fluid flow problems*, J. Comput. Phys., 130 (1997), pp. 41–53.
- [16] A. OVERMAN AND J. V. ROSENDALE, *Mapping robust parallel multigrid algorithms to scalable memory architectures*, in Proceedings of the Sixth Copper Mountain Conference on Multigrid Methods, 1993.
- [17] K. STÜBEN, *Parallel and robust multigrid with applications*, in Proceedings of the Euroconference on Supercomputation in Nonlinear and Disordered Systems: Algorithms, Applications and Architectures, L. Vázquez, F. Tirado, and I. M. Llorente, eds., World Scientific, 1997.
- [18] C. THOLE AND U. TROTTENBERG, *Basic smoothing procedures for the multigrid treatment of elliptic 3d operators*, Appl. Math. and Comput., 19 (1986), pp. 333–345.
- [19] T. WASHIO AND K. OOSTERLEE, *An evaluation of parallel multigrid as a solver and a preconditioner for singularly perturbed problems, Part II: Flexible 2D and 3D multiple semicoarsening*, Tech. Report 1012, GMD, 1996, to appear in SIAM J. Sci. Comp.
- [20] P. WESSELING, *An Introduction to Multigrid Methods*, John Wiley & Sons, New York, 1992.
- [21] I. YAVNEH, *Multigrid smoothing factors of red-black Gauss-Seidel applied to a class of elliptic operators*, SIAM J. Numer. Anal., 32 (1995), pp. 1126–1138.

| REPORT DOCUMENTATION PAGE | | | Form Approved OMB No. 0704-0188 | |
|--|--|---|---|--|
| Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503. | | | | |
| 1. AGENCY USE ONLY (Leave blank) | 2. REPORT DATE May 1999 | 3. REPORT TYPE AND DATES COVERED Contractor Report | | |
| 4. TITLE AND SUBTITLE Plane smoothers for multiblock grids: Computational aspects | | | 5. FUNDING NUMBERS C NAS1-97046 WU 505-90-52-01 | |
| 6. AUTHOR(S) Ignacio M. Llorente Boris Diskin N. Duane Melson | | | | |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Computer Applications in Science and Engineering Mail Stop 132C, NASA Langley Research Center Hampton, VA 23681-2199 | | | 8. PERFORMING ORGANIZATION REPORT NUMBER ICASE Report No. 99-17 | |
| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Langley Research Center Hampton, VA 23681-2199 | | | 10. SPONSORING/MONITORING AGENCY REPORT NUMBER NASA/CR-1999-209331 ICASE Report No. 99-17 | |
| 11. SUPPLEMENTARY NOTES Langley Technical Monitor: Dennis M. Bushnell Final Report To be submitted to Journal of Computational Physics | | | | |
| 12a. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified-Unlimited Subject Category 64 Distribution: Nonstandard Availability: NASA-CASI (301) 621-0390 | | | 12b. DISTRIBUTION CODE | |
| 13. ABSTRACT (Maximum 200 words) Standard multigrid methods are not well suited for problems with anisotropic discrete operators, which can occur, for example, on grids that are stretched in order to resolve a boundary layer. One of the most efficient approaches to yield robust methods is the combination of standard coarsening with alternating-direction plane relaxation in the three dimensions. However, this approach may be difficult to implement in codes with multiblock structured grids because there may be no natural definition of global lines or planes. This inherent obstacle limits the range of an implicit smoother to only the portion of the computational domain in the current block. This report studies in detail, both numerically and analytically, the behavior of blockwise plane smoothers in order to provide guidance to engineers who use block-structured grids. The results obtained so far show alternating-direction plane smoothers to be very robust, even on multiblock grids. In common computational fluid dynamics multiblock simulations, where the number of subdomains crossed by the line of a strong anisotropy is low (up to four), textbook multigrid convergence rates can be obtained with a small overlap of cells between neighboring blocks. | | | | |
| 14. SUBJECT TERMS robust multigrid methods, multiblock grids, anisotropic discrete operators | | | 15. NUMBER OF PAGES 29 | |
| | | | 16. PRICE CODE A03 | |
| 17. SECURITY CLASSIFICATION OF REPORT Unclassified | 18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified | 19. SECURITY CLASSIFICATION OF ABSTRACT | 20. LIMITATION OF ABSTRACT | |

